1

Making Ad Hoc Networks Density Adaptive

Ram Ramanathan
Internetwork Research Department
BBN Technologies (A part of Verizon)
Cambridge, Massachusetts, USA
ramanath@bbn.com

Abstract- Many ad hoc networks, including military networks, have a wide range of deployment densities. Further, they exhibit mobility patterns that induce continual and rapid density changes (for instance, due to the confluence and dispersion of troops). Most existing ad hoc networking protocols tacitly or explicitly ignore density variations. In this paper, we discuss the performance degradation caused by density changes, and describe techniques for resilience to density variations. In particular, we study the impact of transmit power control and directional antennas on network performance and show that significant improvements can be obtained.

I. INTRODUCTION

Density is the number of network nodes per unit area. For a given transmission range, as density increases, each node gets a smaller fraction of the channel for itself. On the other hand, the average number of hops for each packet decreases. The decrease in path length, however, is linear in the transmission range, whereas the increase in number of nodes contending is quadratic in the range, and so one expects that an increase in density would result in reduced capacity after a certain point. This hypothesis has been validated by a number of theoretical and experimental work including [1], [2], [3]. For instance, in [2] it has been shown that the smallest node degree that results in a connected network maximizes the throughput, and in [1] it is shown that fairly small node degree (around 8) would serve throughput the best.

While high density is bad for effective capacity, low density is also bad, but for a different reason – *connectivity.* For a given transmission range, as density decreases, the connectivity redundancy of the network decreases, and the probability of partitioning increases. Thus, protocols are needed that work under low as well as high densities, and for mobile networks, adapt to changes in density.

We consider two ways of making ad hoc networks density adaptive – by using transmission range (power) control, and by using beamforming antennas. The idea behind power control is that by reducing the transmission power in accordance with the density, one can control the node degree to be within a desirable range. The idea behind beamforming antennas is that they provide automatic spatial reuse (due to smaller beamwidth) which helps at high densities, and offer longer ranges (due to increased gains) which helps at lower densities. To exploit power control and beamforming antennas, however, new network layer algorithms are required. This is the subject of this paper.

The power control idea has thus far been explored in two directions. The first is to use *topology control* [3], [4]. That is, the physical topology is not taken as a given. Rather, an attempt is made to (adaptively) mould the topology so as to maintain a given objective, for instance: *within the constraint of keeping the network connected, minimize the node degree.* Specifically, the problem is one of dynamic assignment of power levels to nodes so as to meet the topology objective. This is a "global", network-layer solution.

The second direction is the use of a per-packet power control at a MAC layer that uses an RTS/CTS or similar handshake before sending a packet. The presence of this handshake allows one to determine if the transmission power is in excess of what is strictly required. In order to reduce interference, the receiver may include in the CTS, the exact power that the sender must use to send that packet. This approach is used by military radios that can do per-packet power control (e.g. the radio used in the DARPA RAVEN [5] project). This is a "local", link-layer solution.

Each of these approaches have their advantages and disadvantages. In this paper, we consider the topology control approach, as it is compatible with a variety of MAC protocols, does not require potentially time consuming power changes in the forwarding path, and does not require per-packet power control. We build upon the work in [3] by introducing more algorithms and providing decentralized versions of algorithms in [3].

The beamforming idea has not been explored very much. Our study of this is restricted to evaluating the performance gains that such antennas can provide even with simplistic algorithms, rather than inventing new algorithms.

II. DENSITY ADAPTATION USING TOPOLOGY CONTROL

We now describe a mechanism called GIFT (Global Information Full Topology). As the name indicates, global information is used to choose a topology from amongst all possible topologies. The objective of GIFT is to minimize a user-specified cost metric under a user-specified connectivity constraint. Cost metrics may be any one of *average power*, *maximum power*, and *maximum degree.* Connectivity constraints specifiable are *connected* network and *biconnected* network. A network is biconnected if and only if the loss of any one node does not partition the network. We first describe the algorithms in a centralized manner and then describe two decentralized versions.

A. GIFT Algorithms

A GIFT algorithm produces an assignment of transmit powers to nodes based on certain minimization objectives and connectivity constraints. We have considered the following objectives and constraints:

- *Minimization objectives.* Maximum power, average power, maximum degree.
- *Constraints.* Connectivity, Biconnectivity.

There are 6 combinations possible. We have developed algorithms for 4 of those combinations. We describe the operation of these algorithms below.

Partially supported by the DARPA Global Mobile Information Systems (GloMo) contract DAAB07-97-C-D606

The algorithms are based on the concept of *Link Closure Power* (LCP). The Link Closure Power for a link between nodes s_1 and s_2, or LCP(s_1, s_2) is the minimum amount of power required to create a link between the nodes. The input to each algorithm is the set of nodes comprising the ad hoc network and the link closure power between every pair of nodes. The output is an assignment of power level for each node.

We have considered only bidirectional links for GIFT, although most of this work is easily extendable for topology construction with unidirectional links.

The first two algorithms presented below were also given in [3]. However, no decentralization was given and hence they were only applicable to stationary networks. Here we describe, in section II-B, a decentralization scheme that is applicable to all of the algorithms.

A.1 Connected Minimum Maximum Power (C-MMP)

Algorithm C-MMP is given formally in the box below. It is a simple "greedy" algorithm, similar to the minimum cost spanning tree algorithm. It works by iteratively merging connected components until there is just one. Initially, each node is its own component. Node pairs are selected in non-decreasing order of their LCPs. If the nodes s_1 and s_2 are in different components, then the transmit power of each is increased to be equal to LCP(s_1, s_2). This is done until the network is connected. The description assumes for simplicity that network connectivity can be achieved without exceeding the maximum possible transmission powers. However, the algorithm can be easily modified to return a failure indication if this is not true.

We prove, in [3], that this is an optimal algorithm, that is, the resulting maximum power is the minimum possible. The running time of the algorithm is $O(n^2 \cdot log(n))$.

Algorithm C-MMP

```
begin
sort node pairs in non-decreasing order
  of LCP
initialize $|N|$ clusters, one per node
for each (u,v) in sorted order do
    if cluster(u) $\neq$ cluster(v)
        $p(u) = p(v) = LCP(u,v)$
        merge cluster(u) with cluster(v)
        if number of clusters is 1
        then end
end
```

A.2 Biconnected Minimum Maximum Power (B-MMP)

First C-MMP is executed to make the network connected. Then, it is augmented to a biconnected network, as described below.

We first identify the biconnected components in the graph induced by the power assignment from algorithm C-MMP. This is done using a standard method based on depth-first search given in [6]. Then, node pairs are selected in non-decreasing order of their mutual distance and joined only if they are in different biconnected components. This is continued until the network is biconnected.

We prove, in [3], that this is an optimal algorithm, that is, the resulting maximum power is the minimum possible. The running time of the algorithm is $O(n^2 \cdot log(n))$.

Algorithm B-MMP

```
begin
execute C-MMP to produce G
for each $(u,v)$ in sorted order do
    if biconn-comp(G,u) $\neq$ biconn-comp(G,v)
        $q$ = LCP(u,v)
        $p(u)$ = max($q$, $p(u)$)
        $p(v)$ = max($q$, $p(v)$)
        add $(u,v)$ to G
end
```

A.3 Connected Minimum Average Power (C-MAP)

The main difference between this and C-MMP is that rather than the LCP, the *incremental* LCP is used. We first define this notion and then describe the algorithm.

The incremental LCP is the sum of the power increases that two nodes would have to make, over and above the assigned power, in order to close the link between them. Formally,

$$incLCP(u,v) \;=\; incLCP(u \rightarrow v) + incLCP(v \rightarrow u) \qquad (1)$$

where

$$incLCP(x \rightarrow y) = MAX(0,(LCP(x,y) - p(x))) \qquad (2)$$

where $x \rightarrow y$ indicates a link from x to y, and $p(x)$ is the power currently assigned to node x.

It works by iteratively merging connected components until there is just one. Initially, each node is its own component, and has a power assignment of zero. At each step in the iteration, we determine the node pair (link) with the minimum incremental LCP (ties broken randomly).

At each step, since $LCP(u,v)$, $p(u)$, and $p(v)$ are known, one can compute $incLCP$. The link between node pairs with the minimum incremental LCP is then closed by assigning the power $LCP(u,v)$ to each node.

This algorithm is not an optimal one. Indeed, the problem has been proven to be NP-complete [7], which means that it is extremely unlikely that there exists a polynomial time optimal algorithm.

Algorithm C-MAP

```
begin
sort node pairs in non-decreasing order
  of LCP
initialize $|N|$ clusters, one per node
while there are more than 1 cluster do
    find $u$, $v$ in different clusters such
    that $incLCP(u,v)$ is a minimum
        $p(u)$ = MAX($p(u)$, $LCP(u,v)$)
        $p(v)$ = MAX($p(v)$, $LCP(u,v)$)
            merge cluster(u) with cluster(v)
end
```

A.4 Connected Minimum Maximum Degree (C-MMD)

C-MMD works by iteratively merging connected components until there is just one. Initially, each node is its own component, and has a power assignment of zero. At each step in the iteration, we determine which unclosed link, if closed, will result in the minimum maximum degree (ties broken randomly). We then increase the power of the relevant nodes to close this link.

```
Algorithm C-MMD
begin
sort node pairs in non-decreasing order
  of LCP
initialize $|N|$ clusters, one per node
while there are more than 1 cluster do
    find $u$, $v$ in different clusters such
    that $G$ with $(u, v)$ has a minimum max-degree
        $p(u)$ = MAX($p(u)$, $LCP(u, v)$)
        $p(v)$ = MAX($p(v)$, $LCP(u, v)$)
            merge cluster(u) with cluster(v)
end
```

B. Decentralized GIFT

Unlike a "fully distributed" algorithm, in decentralized GIFT all nodes have the same information through a flooded exchange, and run the same algorithm to produce the same result, after which each node uses the part of the result that pertains to itself. This is similar to the way the traditional link-state protocol works.

We have designed two variants of GIFT:

- Position-based GIFT (P-GIFT). This uses geographical position information, derived from GPS or any other source, in order to compute the power assignments.
- Signal Strength based GIFT (S-GIFT). This does not use any position information. Instead, it uses the Received Signal Strength Information (RSSI) from each radio, along with a transmitted power value, in order to compute the power assignments.

Both P-GIFT and S-GIFT use some (the same) topology control algorithms described in section II-A that form the "kernel" of these mechanisms. A main difference between S-GIFT and P-GIFT is the way they derive the values for LCP. We now describe P-GIFT and S-GIFT in detail.

B.1 Position-based GIFT (P-GIFT)

The execution of P-GIFT can be perceived as "punctuated equilibrium", with things happening at globally synchronized periodic intervals (for instance, every minute, on the minute). The times at which these punctuations happen are called Topology Reformation Moments (TRM). At each TRM, the following things happen.

First, each node triggers a position update to all other nodes using a flooding protocol (exactly the same as used for link state updates). The position update contains the current GPS position of the node. In order to avoid collisions with neighboring nodes' position updates, the actual sending of the position update is jittered randomly.

Next, each node waits for a specified period of time to collect updates from nodes. This period is large enough to allow for position updates to propagate thoughout the network (higher for a bigger/sparser network), but much smaller than the interval between TRMs.

At the end of this period, each node runs one of the topology control algorithms specified in section II-A (user choice), and computes the tranmsit power of each node. The *link closure power* specified as input to those algorithms is determined as a function of the distance between the nodes forming the link. The exact function should be chosen based on the path loss characteristics of the terrain in question. We have used the following function, based on [9]:

$$LCP(u, v) = \Lambda(d) + \tau \tag{3}$$

where d is the euclidean distance between u and v, τ is the threshold receiver sensitivity, and Λ is given by

$$\Lambda(d) = \Lambda_{ref} + 10 \cdot \mathcal{E} \cdot log_{10}(\frac{d}{d_{ref}}) \tag{4}$$

where d_{ref} is a reference distance and Λ_{ref} is the attenuation at the reference distance. These values can be found by empirical measurement or approximately as a function of wavelength and antenna dimensions and gains, as given in [9]. We have done the latter.

Since the same position information is used by each node (barring lost packets), the same result is obtained by each node.

Finally, a node determines if in the resulting power assignment, is different from its current operational power, and if so adjusts it. The position information is then discarded and the nodes use this operational power until the next TRM.

B.2 Signal strength based GIFT (S-GIFT)

There are two shortcomings in P-GIFT. First, it needs position information through GPS or other means which adds complexity to the node, and may not be available indoors. Second, the LCP uses an "estimated" path loss function, which is at best an approximation to the real path loss, and indeed may be a very bad approximation when there are obstacles.

In contrast, S-GIFT does not use position information at all. The basic idea is that every node determines the LCP for the link from each of its neighbors to itself by using the value of the transmitted power and the received power. The Received Signal Strength Indication (RSSI) is available for many commercial radios, including the radio we used (see section III-A for details on the radio).

The execution of S-GIFT is also decentralized in the same manner as P-GIFT, in particular, as Topology Reformation Moments. At each TRM, however, different things happen, as follows.

Each node broadcasts a Power Inclusive Probe (PIP) that contains its identifier and the power P at which the PIP is sent (by default the maximum possible power). A node N that receives a PIP from a node M notes the received signal strength S. It also has a receive threshold T, which is the minimum signal strength it needs to receive the packet correctly. If S is less than T, then the packet cannot be received. So if the packet is received, we assume that S is greater than T. The value of S - T denotes the "excess" power. In other words, it was sufficient for the transmitter to have transmitted at P - (S - T) for the packet to have been received correctly. Node N stores this value as the LCP between M and N. Assuming symmetric path losses and the same T, this should be equal to the value computed by M for the same link.

Once the LCP for all neighbors is calculated, a node floods this information in an LCPInfo message throughout the network as part of the decentralized process (similar to the flooding of position information in P-GIFT). At the end of the specified period for collecting updates, each node is in a position to run one of the algorithms described in section II-A since it knows the LCP between every pair of nodes. Since the same information is used by each node (barring lost packets), and the same algorithm is run, the same power assignments are computed at each node. A node changes its operational power if the computation says that its own new power is different.

B.3 Overhead

The average control overhead for P-GIFT and S-GIFT is given by $N \cdot K/T$, where N is the number of nodes, K is the control message size and T is the inter-TRM interval. Using T, one can tradeoff the overhead versus responsiveness. In our simulations, we used a value of 20 seconds for T, which was found to be adequate. With this, and assuming K conservatively as 20 bytes, the per node overhead for a 1000 node network is only about 9 kbps, which is less than 0.5 % of a 2 Mbps transceiver. Thus, this is a simple yet scalable approach for most ad hoc networks.

C. Experimental Results

The experimental results are based on a very high fidelity simulation using C++ of P-GIFT in the context of an ad hoc networking system that uses a flat *link-state routing* mechanism. Event driven and periodic updates are flooded throughout the network and routes are generated using Dijkstra's shortest path algorithm. For details in this system, please refer [8].

A detailed model of the Utilicom Longranger 2050 [10] radio is used. It is a direct sequence spread spectrum radio in the 2.4 GHz ISM band capable of a raw data rate of about 1676 Kbps. The Utilicom radio has transmit power control, and is the basic reason for our choice. It uses a proprietary collision avoidance protocol for channel access.

We use a psuedo-random mobility model and an r^4 propagation model. Each node does an SNIR calculation to determine whether to mark the packet as successful or collided. The traffic load consists of 30 streams between randomly chosen endpoints, each stream consisting of uniformly distributed packets at 3.9 packets/sec. Each packet is 256 bytes, resulting in 8 kbps streams. The interval between Topology Reformation Moments (TRMs) was set to 20 seconds.

In results not presented here, we found that B-MMP outperformed C-MMP, C-MMD, and C-MAP for stationary networks. We expect that superiority to hold for mobile networks as well, and therefore only study B-MMP further here.

The decentralized B-MMP GIFT is compared with the Local Information No Topology (LINT) algorithm presented in [3] and to the performance when no topology control is employed. In LINT, each node periodically adjusts its power so as to keep its node degree (number of neighbors) within a high and low threshold. Power is reduced (increased) using a formula (refer [3]) if the degree is higher (lower) the threshold. All results reported here are for 60 nodes.

We note from figure 1 that the GIFT algorithm yields the best throughput amongst all mechanisms studied. The throughput at density 14 is about 11% better than that of LINT, and about 71% better than no topology control. However, the delay is also slightly higher, about 27% more than LINT. The increased delay is a result of being less densely connected (giving rise to larger number of hops between node pairs). Although this should reduce the MAC delay, we are not sure that the proprietary MAC protocol is exploiting that adequately. Both throughput and delay of GIFT are largely invariant over a wide range of densities, attesting to its density-adaptivity.

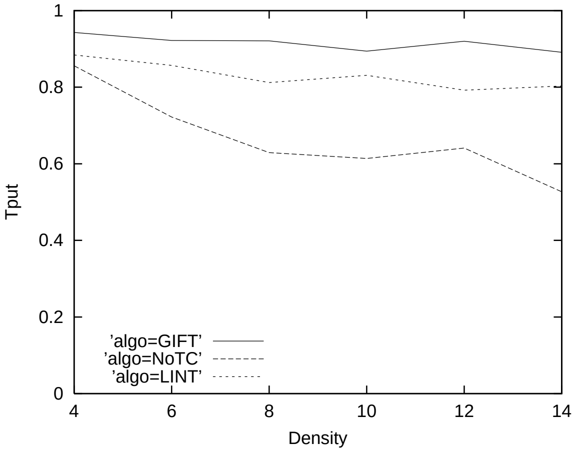

Tput vs Density for various nbrMax; 60 nodes; Mobility=0.001, VC=N, LST=T

Fig. 1. Throughput vs Density

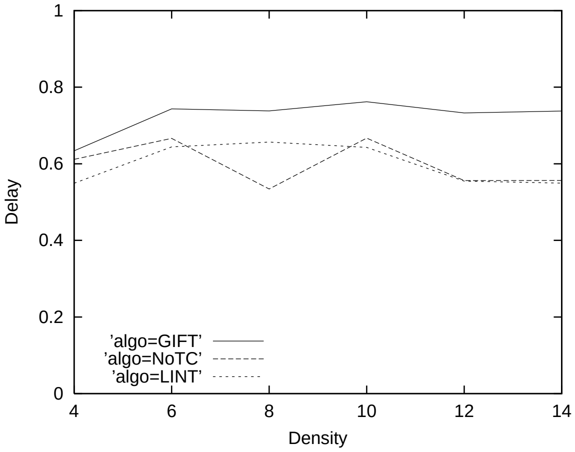

Delay vs Density for various nbrMax; 60 nodes; Mobility=0.001, VC=N, LST=T

Fig. 2. Delay vs Density

III. DENSITY ADAPTATION USING BEAMFORMING ANTENNAS

This part of the study was conducted using the OPNET modeling and simulation tool (version 6.0) as it provided the the best support for modeling directional communications. The neighbor discovery and link-state based routing described in the previous section were modeled. In addition, beamforming was introduced as described below.

Each node in the simulation system has an antenna that can be associated with one of a number of predefined antenna patterns. Beam steering is modeled by orienting the pattern so that the center of the main lobe points toward the target node. OPNET allows both the antenna pattern associated with a node, as well as the orientation of the pattern to be changed dynamically (on a per packet basis), and this allows us to model both steerable and switched beam antennas.

Beam steering is modeled as follows. Assume that a pattern P_g corresponding to gain g has been associated with a node S. Suppose that S wants to send a packet using a steerable beam to another node R. After S has obtained access to the channel, it sets its pattern to P_g and the main beam is pointed in the direction of R. The packet is then sent for transmission. Immediately after the packet is transmitted, S sets its pattern back to "OMNI". Thus, we only use transmit directionality, all receives are omni-directional.

The channel access is changed to an RTS/CTS based protocol, but a very aggressive one. In particular, we model a protocol in which a node is *never* blocked upon receiving an

RTS or CTS. The handshake is used only for ensuring that the receiver is not already busy sending or receiving. In particular, no attention is paid to the receive status of other nodes. Thus, this could potentially cause collision at nodes other than the intended receiver. However, with the fairly narrow beams, this tends to seldom happen.

A. *Experimental Results*

The throughput (see figure 3) with beamforming antennas is significantly higher than with omni-directional antennas. Interestingly, the greatest difference occurs at middle densities – e.g., at density of 48, using a 26 dBi antenna gives 28% better throughput than the omni directional antenna. This is a result of counteracting forces – the shorter number of hops, whose beneficial effects increase with increasing density, versus the detrimental effect of sidelobes, which also increases with increasing density, resulting in a peak at the middle densities. The difference between the various antenna gains is much less - to within 6% in most cases, and about 10% in the density 16 case.

The delay (see figure 4) for steered beams with aggressive CA and power control is dramatically lower than with omni-directional antennas (also using aggressive CA and power control). When the density is 112 nodes/sq mile, there is a reduction by a factor of about 28 in the delay when the 26 dBi antenna is used. The difference is less at lower densities (about a factor of 2-5 at density 16). The delay of both omni-directional as well as beamforming antennas is higher at lower densities. This is due mainly to the increased number of hops at lower densities.

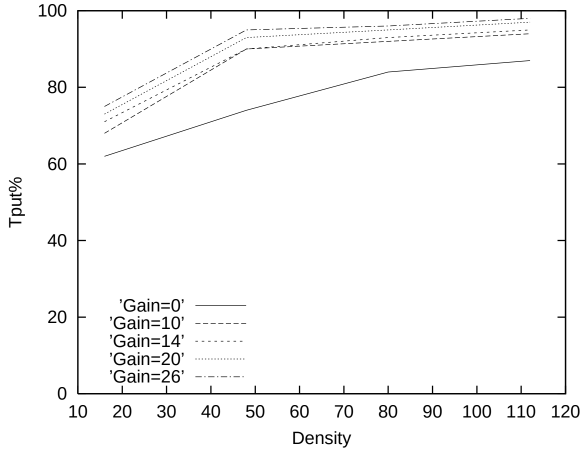

Tput% vs Density for various Gain; 40 nodes; NumAnt=1

Fig. 3. Throughput vs density, gains due to spatial reuse only, steerable antennas, static networks

IV. Concluding Remarks

Density adaptation is an important part of ad hoc network design. We have described two distinct and orthogonal ways of making ad hoc networks density adaptive, namely, using topology control and steerable antennas.

For topology control, we have built upon previous work and presented new and decentralized algorithms for "global" topology control. Our experimental results indicate that the new protocol improves performance by about 71 %, and outperforms previously proposed LINT by about 11 %.

We have also shown using comprehensive simulations that the use of beamforming antennas can greatly improve capacity just due to the increased spatial reuse. In particular, even using a simplistic MAC protocol, using steerable antennas results in up to 28 % improvement in throughput and up to a factor of 28 reduction in delay for a 26 dBi antenna.

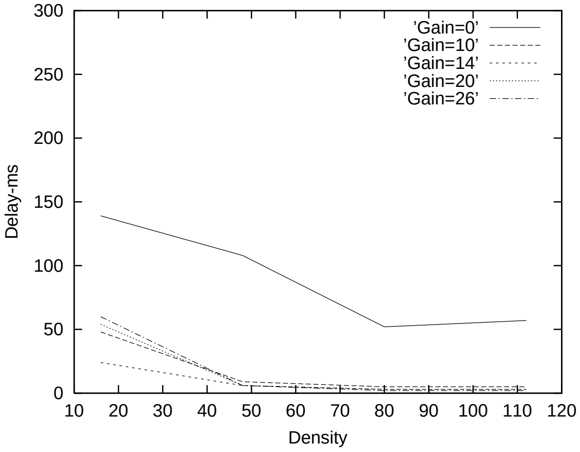

Delay-ms vs Density for various Gain; 40 nodes; NumAnt=1

Fig. 4. Delay vs density, gains due to spatial reuse only, steerable antennas, static networks

In sum, our results show that density adaptation can provide a significant performance gain. However, this is just the tip of the iceberg, and more work needs to be done to study density adaptation

References

[1] H. Takagi, L. Kleinrock, "Optimal transmission ranges for randomly distributed packet radio terminals", *IEEE Transactions on Communications*, COM-32, No. 3, March 1984.

[2] P. Gupta and P.R. Kumar, "The Capacity of Wireless Networks", *IEEE Transactions on Information Theory*, vol. IT-46, no. 2, pp. 388-404, March 2000.

[3] R. Ramanathan and R. Hain, "Topology Control of Multihop Radio Networks using Transmit Power Adjustment," in *Proceedings of IEEE Infocom*, Tel Aviv, Israel, 2000.

[4] L. Hu, "Topology Control For Multihop Packet Radio Networks", Proc. IEEE Infocom, 1991.

[5] Reliable All-informed Voice Networks - RAVEN - Final Technical Report, available from CECOM, reference DARPA contract DAAB07-97-C-D012.

[6] R. Sedgewick, *Algorithms*, Addison-Wesley, 1983.

[7] Clementi, Penna, Silvestri, "Power Range Assignment Problem in Packet Radio Networks", in *Proceedings of APPROX 99*, Berkeley, CA.

[8] S. Ramanathan, M. Steenstrup, "Hierarchically-organized, Multihop Mobile Networks for Multimedia Support", *ACM/Baltzer Mobile Networks and Applications*, Vol. 3, No. 1, pp 101-119.

[9] T. S. Rappaport, *Wireless Communications, Principles and Practice*, Prentice-Hall, 1996.

[10] www.utilicom.com